Journal of Physics: Conference Series

Data Center Environmental Sensor for safeguarding the CERN data archive

To cite this article: J Leduc *et al* 2017 *J. Phys.: Conf. Ser.* **898** 062053

Data Center Environmental Sensor for safeguarding the CERN data archive

J Leduc[1], V Bahyl[1], G Cancio[1], E Cano[1], Q Genoud[2], D F Kruse[1], and S Murray[1]

[1]European Organization for Nuclear Research CERN, CH-1211 Genève 23
[2]GEII Université de Savoie, Annecy le Vieux, France

E-mail: {`Julien.Leduc, Vladimir.Bahyl, German.Cancio.Melia, Eric.Cano, Daniele.Francesco.Kruse, Steven Murray`}`@cern.ch`

Abstract. CERN has been archiving data on tapes in its Computer Center for decades and its archive system is now holding more than 135 PB of HEP data in its premises on high density tapes. For the last 20 years, tape areal bit density has been doubling every 30 months, closely following HEP data growth trends. During this period, bits on the tape magnetic substrate have been shrinking exponentially; today's bits are now smaller than most airborne dust particles or even bacteria. Therefore tape media is now more sensitive to contamination from airborne dust particles that can land on the rollers, reels or heads. These can cause scratches on the tape media as it is being mounted or wound on the tape drive resulting in the loss of significant amounts of data. To mitigate this threat, CERN has prototyped and built custom environmental sensors that are hosted in the production tape libraries, sampling the same airflow as the surrounding drives. This paper will expose the problems and challenges we are facing and the solutions we developed in production to better monitor CERN Computer Center environment in tape libraries and to limit the impact of airborne particles on the LHC data.

1. Why develop a Data Center Environmental Sensor (*DCES*)?

The physical data bits of today's tapes are smaller than most airborne dust particles or even bacteria. One big data preservation issue comes from the fact that those precious bits are directly exposed to the outside world during read and write operations.

Every time a tape is mounted in a tape drive, the tape medium is wound around reels inside the tape drives, directly exposing the bits to the traversing drive airflow. Any airborne particles passing through the drive can be trapped between two layers of tape medium while winding the medium around reels.

The trapped particles are in direct contact with the magnetic side of the mounted tape (see Figure 1).

To make things worse the tape medium travels at high velocity within a tape drive. This means airborne particles can smash into the magnetic substrate with pressure that can cause severe physical damage to the tape. Since the bits are much smaller than those airborne particles, each impact destroys several bits of data, sometimes beyond ECC correction capabilities.

With current tape library and drive architectures, dust pollution induced data loss will only become a more prevalent problem. The physical size of data bits on tape is decreasing with each

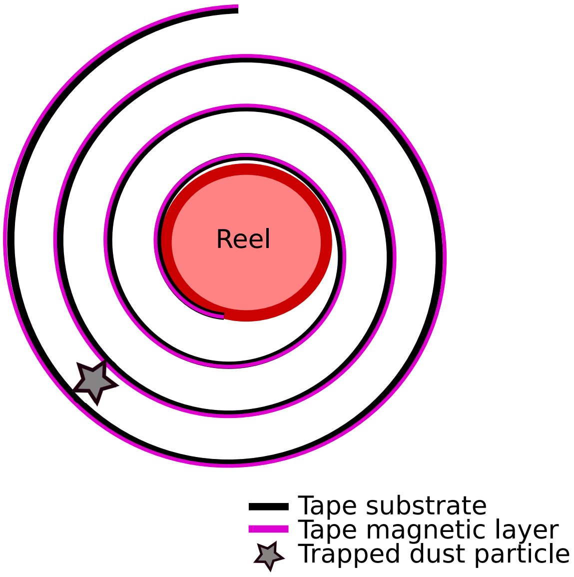

Figure 1. Trapped dust particle damaging tape magnetic layer.

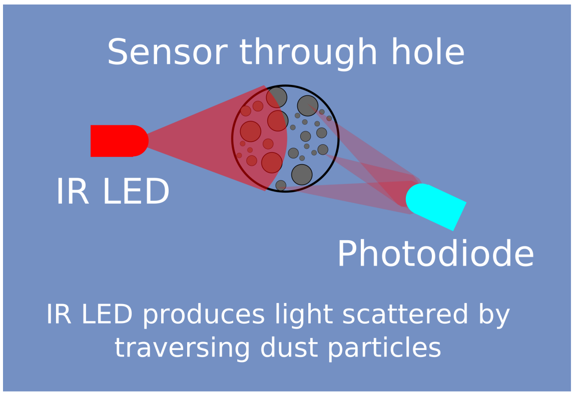

Figure 2. Dust sensor internals.

new generation of tape. This means future tapes will be more sensitive to smaller and smaller dust particles.

The definitive solution to this problem is to use tape library designs that filter or isolate outside air in order to protect the tape medium. Unfortunately the tape libraries at CERN do not have such filters. As tape libraries are a long term investment, the libraries used at CERN still have several years of expected use.

The situation is worsened by the fact that the tape libraries are airflow highways. Each Oracle SL8500 library absorbs more than $33m^3$ of air per minute through its front panel. This massive airflow collects all of the airborne particles in suspension around the library and distributes them across all the tape drives within the library. Each drive takes its share of the inlet airflow which is approximately $0.6m^3$ of air per minute.

In addition to the possible impacts of dust pollution on tape stored data, tape libraries have environmental metric ranges that need to be met (Table 1 gives an example). These metrics must be monitored and the environment modified accordingly in order to ensure environment conformity. Indeed, in modern computer center environments, free cooling can significantly increase relative humidity, if used during warm summer nights, or significantly lower it, during cold dry winter days.

To mitigate the threat of dust to its tape libraries, CERN has prototyped and built custom environmental sensors hosted inside the production tape libraries that sample the same airflow seen by the tape drives.

1.1. Available turnkey solutions

There are two main categories of commercially available turnkey solutions, those of clean room monitoring and those of environmental pollution monitoring in polluted areas of cities.

Clean rooms are specifically designed to minimize dust propagation and sensors are sized for full room monitoring. These solutions require the physical installation of sampling pipes, dust sensors, vacuum pumps and multiplexers. All of which require regular mechanical maintenance (pump and multiplexer), cleaning (sampling pipes and sensors) and calibration (sensors). All those constraints mean that room-scale real-time dust monitoring is prohibitively expensive for

Table 1. Oracle SL8500 environmental specifications during operation.

Environmental	Measurement range
Temperature	+16°C to +32°C
Relative Humidity	20% to 80%
Airborne particles	ISO 14644-1 Class 8

both acquisition and maintenance. Another limitation of this solution is that it relies on time based multiplexing and it can simply miss short lived airborne particle pollution at a single location.

The sensors used for environmental pollution monitoring measure slowly evolving metrics and provide a stable reading every 30 seconds. This rate is far too slow for tape drives. A tape drive can read or write 30GB of data in 30 seconds and it can move a lot more through the drive if it is positioning the tape and not transferring data. Another limitation with environmental pollution sensors is that they sample air at rest and react poorly to high speed airflows which quickly obstruct the sensors with dust. One such sensor was tested at CERN; it read measurements for 2 months and then completely stopped working displaying only zeros.

The two turnkey solutions mentioned above do not physically fit inside a free drive slot within a tape library. The solutions would therefore require the sensors to be installed around the tape libraries. The sensors would take indirect measurements of the possible dust going through the tape drives. An airflow model would have to be used in order to deduce the actual dust levels in the tape libraries from the indirect measurements taken by the sensors.

1.2. DCES: CERN prototype sensor

As turnkey solutions are too expensive and impractical, a prototype sensor had to be built to fulfill our specific requirements. The tape libraries and their vendors impose form factor and other warranty related constraints on the dust sensor. All the components were installed in a decommissioned tape drive tray, allowing the sensor to be installed in a tape library drive slot. The vendor was contacted and allowed CERN to proceed with the sensor installation in production libraries without voiding the warranty. This step was essential before proceeding further as voiding warranty would have been a showstopper for this project.

A kind reminder and disclaimer: anyone willing to develop his own sensor has the responsibility to individually contact the vendor of the equipment hosting it. Neither the authors, nor CERN, can be held responsible of any consequence resulting from unauthorized sensor installation.

Details regarding the design will be discussed in later sections.

2. How does it work?

The DCES is composed of 3 types of sensor: relative humidity, temperature and dust. There are plenty of precise, cheap and calibrated electronic sensors to measure relative humidity and temperature. The same is not true for dust.

The `SHARP GP2Y1010AU0F`[1] sensor is very cheap, robust and widely used in harsh environments (industrial Heating, Ventilation and Air Conditioning installations for example). Like many dust sensors, this is an optical sensor that relies on optical diffraction intensity measurements to evaluate dust particle density.

Figure 2 illustrates how the sensor works. Dust particles travel through a hole in the sensor.

This hole is illuminated by an infrared LED and the density of the particles is derived from the diffracted light intensity. A photodiode on the other side of the sensor measures the scattered light intensity. An Arduino microcontroller samples these readings at 100Hz and computes the average and standard deviation of the 100 measurements performed during the past second.

2.1. Small dust particle density measurements

Turnkey environmental dust sensors need around thirty seconds to acquire enough data to compute an accurate measurement. The DCES microcode implements several tricks to lower that measurement period to only 1.5 seconds, making it 20 times faster.

But there is still one major problem remaining: the off the shelf SHARP sensor is not calibrated. Two approaches to the calibration issue are:

(i) calibrate the sensor.
(ii) completely skip it and work with uncalibrated data.

At CERN a DCES and a calibrated dust sensor were installed together in a sealed box and volatile chemicals sprayed inside. The two sensors measured the same airflow containing the same density of volatile chemical droplets.

Both readings were collected and linear calibration factors determined. Figure 3 shows the result from both sensors with the linear calibration factors applied to the DCES.

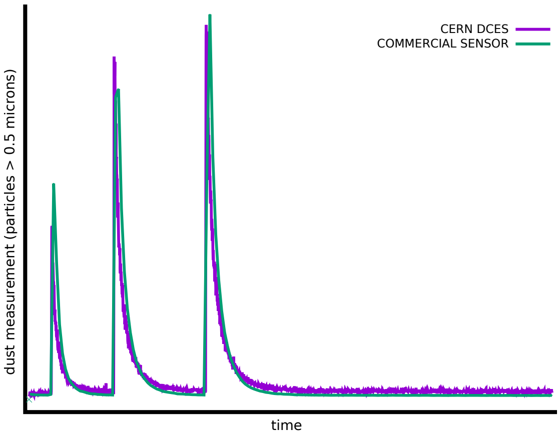

Figure 3. Commercial dust sensor compared to calibrated DCES sensor.

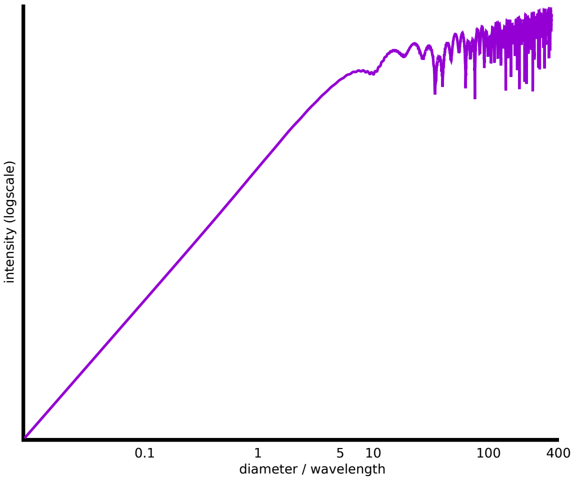

Figure 4. Measured intensity according to size factor.

According to Figure 3 the calibrated DCES sensor is 20 times faster and more precise than the turnkey sensor: **it could detect small dust particle peaks earlier and closer to their highest intensity than the commercial dust sensor**.

This illustrates that fast measurements are very important when measuring dust densities as dust levels **decrease exponentially over time**.

As said earlier, the other possibility is to completely skip calibration. For most of the time, a data center environment is stable and clean. Excessive dust pollution is an abnormal event. This assumption completely removes the need for calibration. In practice CERN environmental monitoring is performed on uncalibrated readings as strict calibration requires regular downtime for the sensors and relaxed calibration adjustments can be applied on the measurements a posteriori.

Experience has shown that **measurements of small dust particles are only meaningful at low speed airflows of a few liters of air per minute**.

2.2. Large dust particle density measurements

The simple and robust `SHARP GP2Y1010AU0F` dust sensor is a single wavelength sensor and therefore cannot directly measure the density of large dust particles. This limitation can be mitigated by optical physics. Mie scattering can be applied to this problem because the diameters of dust particles are similar to the wavelength of the infrared LED of the sensor.

Mie theory describes the scattering of light by particles and it can be used to model the light intensity received by the optical sensor photodiode after it has been scattered by an homogeneous cloud of spherical particles of a fixed diameter. The `Mieplot` [2] software was used to simulate the light intensities that the optical sensor would receive for various sizes of dust particles. The resulting graph in Figure 4 illustrates how the measured intensity fluctuates when the diameter of dust particles is over five times bigger than the wavelength of the scattering light (for the SHARP sensor it means particles larger than 4 microns).

The measured average intensity is an almost linearly increasing function of the particle diameter for small dust particles less than 4 microns in diameter. In constant regime we have observed that DCES measured standard deviation remains low: the coefficient of variation - $\sigma/mean$ - stays below 5%.

When dust particles are larger than 4 microns, the measured average intensity is still an increasing function of particle diameter but it is not linear anymore: Mie resonance phenomena appear and the measured light intensity is oscillating. In practice we have observed that, in the presence of large dust particles in the airflow, the DCES measured coefficient of variation jumps over 20%. A common threshold value can be defined on the measured standard deviation for all the DCES sensors: over this threshold an alarm is triggered indicating the presence of large particles in the computer center.

Another nice feature of this criterion is that standard deviation does not depend on the sensor calibration, which means that **no calibration is needed for the DCES dust sensor to accurately qualify the presence of large dust particles in high speed airflows**.

3. Hardened design

A DCES is composed of simple, robust, widely tested, cheap and widely available components. Its brain is an Arduino board that handles real-time Analog to Digital conversions and computations before sending those to an embedded computer.

The `Raspberry Pi 2`[3] was chosen because it is simple, cheap and it does not contain any mechanical components or moving parts.

As stated in earlier sections, the industrial HVAC dust sensor is designed for rough environments and is practically maintenance free in high speed airflows.

As an illustration, the four DCES sensors that have been running inside CERN tape libraries for the last eighteen months did not require any maintenance. They have gone through several power cuts, and have been sustaining stronger airflows than the surrounding tape drives. This allows them to ingest more dust than the tape drives, which means better sampling and additional dust protection for the tapes as the strongest airflow diverts particles from tape drives to DCES sensor.

The DCES has passed tests to qualify it for use in the CERN experiment pits and **has been successfully tested in magnetic fields of 0.5 Teslas** (Figure 5).

4. Easy integration

The CERN DCES hardware design is available as open source[4], under the **CERN Open Hardware License v1.2**[5]. As such, the PCB can be freely modified to fit any form factor, add sensors or external communication interfaces.

Since it is Arduino based there are many external software libraries available that enable fast and easy integration of additional hardware.

Figure 5. DCES electronics during magnetic field torture test.

Figure 6. A DCES retro-fitted inside a standard ATX PSU: portable WIFI DCES.

The portable WIFI version of DCES (Figure 6) illustrates this flexibility: this instance automatically connects to the CERN WIFI network and provides autonomous data logging and monitoring via its embedded web server.

The **DCES custom firmware is also available**[6] **as open source under GPLv3**[7].

A commercial license is available from CERN if GPLv3 is not suitable, please contact us if this applies to you.

The DCES easily integrates into the existing monitoring infrastructure at CERN. Each DCES is Puppet managed and runs a standard SNMP agent that sends SNMP traps when abnormal levels of dust are detected. Dust information is sent through a `lemon` sensor and into CERN `Elasticsearch` instance, and simultaneously into an `InfluxDB` time series database that allows custom analysis in `Jupyter` notebooks.

Integration possibilities are endless thanks to DCES being a fully open source project.

The DCES does not rely on any CERN specific IT infrastructure and has already been or is being integrated outside the CERN IT department.

5. References

[1] `SHARP GP2Y1010AU0F` specifications *https://www.sharpsde.com/products/optoelectronic-components/model/GP2Y1010AU0F*
[2] MiePlot homepage *http://www.philiplaven.com/mieplot.htm*
[3] Raspberry Pi 2 product page *https://www.raspberrypi.org/products/raspberry-pi-2-model-b*
[4] DCES homepage on open hardware repository *http://www.ohwr.org/projects/dces-dtrhf-ser1ch-v1*
[5] The CERN Open Hardware licenses *http://www.ohwr.org/projects/cernohl/wiki*
[6] The DCES firmware repository *http://www.ohwr.org/projects/dces-dtrhf-ser1ch-v1/repository*
[7] The FSF GPLv3 license *https://www.gnu.org/licenses/gpl-3.0.en.html*